Point Process-based Monte Carlo estimation

Clément Walter

Laboratoire de Probabilités et Modèles Aléatoires
Université Paris Diderot, Paris, France
CEA, DAM, DIF, F-91297 Arpajon, France
Tel.: +331-69-264000
E-mail: clement.walter@cea.fr

arXiv:1412.6368v5 [cs.CE] 9 Sep 2015

Abstract This paper addresses the issue of estimating the expectation of a real-valued random variable of the form $X = g(\mathbf{U})$ where g is a deterministic function and $\mathbf{U}$ can be a random finite- or infinite-dimensional vector. Using recent results on rare event simulation, we propose a unified framework for dealing with both probability and mean estimation for such random variables, *i.e.* linking algorithms such as Tootsie Pop Algorithm (TPA) or Last Particle Algorithm with nested sampling. Especially, it extends nested sampling as follows: first the random variable X does not need to be bounded any more: it gives the principle of an ideal estimator with an infinite number of terms that is unbiased and always better than a classical Monte Carlo estimator – in particular it has a finite variance as soon as there exists $k \in \mathbb{R} > 1$ such that $\mathrm{E}\left[X^k\right] < \infty$. Moreover we address the issue of nested sampling termination and show that a random truncation of the sum can preserve unbiasedness while increasing the variance only by a factor up to 2 compared to the ideal case. We also build an unbiased estimator with fixed computational budget which supports a Central Limit Theorem and discuss parallel implementation of nested sampling, which can dramatically reduce its computational cost. Finally we extensively study the case where X is heavy-tailed.

Keywords Nested sampling · Evidence · Central limit theorem · Heavy tails · Trimmed mean · Tail index estimation · Rare event simulation · Last Particle Algorithm

1 Introduction

Nested sampling was introduced in the Bayesian framework by Skilling (2006) as a method for "estimating directly how the likelihood function relates to prior mass". Formally, it builds an approximation for the evidence:

$$Z = \int_\Theta L(\theta)\pi(\theta)\mathrm{d}\theta,$$

where π is the prior distribution, L the likelihood, and $\Theta \subset \mathbb{R}^d$. It is somehow a quadrature formula but in the $[0,1]$ interval rather than in the original multidimensional space Θ:

$$Z = \int_0^1 Q(P)\mathrm{d}P,$$

where Q is the quantile function which is the generalised inverse of:

$$P(\lambda) = \int_{L(\theta)>\lambda} \pi(\theta)\mathrm{d}\theta.$$

Hence the name *nested sampling* because the initial input space is divided into nested subsets $\{\theta \in \Theta \mid L(\theta) > \lambda\}$. Convergence of the approximation error toward a Gaussian distribution has been proved (Chopin and Robert 2010) when assuming that Q is twice continuously differentiable with its two first derivatives bounded over $[\varepsilon, 1]$ for some $\varepsilon > 0$.

On the other hand estimating a quantity such as $P(\lambda)$ for a given λ is a typical problem arising in rare event probability estimation. In this context, L (often denoted by g) represents a complex computer code (not necessarily positive valued nor continuous nor bounded), θ is a vector of parameters, and $F_\lambda = \{\theta \in \Theta \mid L(\theta) > \lambda\}$ is the so-called failure domain. The idea of writing F_λ as a finite intersection

of nested subsets $F_{\lambda_0} \supset \cdots \supset F_{\lambda_n}$, $-\infty = \lambda_0 < \cdots < \lambda_n = \lambda$ goes back to Kahn and Harris (1951) and is now referred to as Multilevel Splitting (Garvels 2000; Cérou and Guyader 2007) or Subset Simulation (Au and Beck 2001). Statistical properties and convergence results have been derived by interpreting the Splitting algorithm in terms of an Interacting Particles System (Cérou et al 2009, 2012). Furthermore a particular implementation, sometimes called the *Last Particle Algorithm* (LPA), has gained a lot of attention and Huber et al (2011, 2014), Guyader et al (2011) and Simonnet (2014) have independently proved its link with a Poisson process. This algorithm is indeed somehow the one proposed by Skilling (2006, Section 6) but the connection between nested sampling and rare event simulation remains unclear (see Guyader et al (2011) and the discussion following Huber et al (2011) in Bernardo et al (2011)).

The goal of this paper is to fill this gap by introducing a common framework for these methodologies. The core tool is that any continuous real-valued random variable can be linked with a Poisson process with parameter 1. Then a family of estimators can be defined using several realisations of such processes instead of *iid* samples. While it only recasts results for extreme probability estimation in a very general setting – *i.e.* the random variable of interest writes as $X = g(\mathbf{U}) \in \mathbb{R}$ where g is a deterministic function and $\mathbf{U}$ can be a random finite- or infinite-dimensional vector – it extends nested sampling to the estimation of the mean of any real-valued random variables (bounded or not) and brings new theoretical results: 1) the ideal estimator with an infinite number of terms (non truncated nested sampling) is unbiased; 2) the ideal nested sampling estimator is always better than the classical Monte Carlo estimator in term of variance; and 3) it has a finite variance as soon as a moment of order $k \in (1,\infty)$ exists.

Moreover we address the issue of the nested sampling termination (see Skilling 2006, Section 7). Using results on Multilevel Monte Carlo (Giles 2008; McLeish 2011; Rhee and Glynn 2013), we show that one can get an unbiased estimator with a random but a.s. finite number of terms whose variance is only twice the one of the ideal estimator. We also build an unbiased estimator with a fixed computational budget which supports a Central Limit Theorem. We further discuss parallel implementation of nested sampling and these new estimators as this can can dramatically reduce its computational cost.

All these theoretical results are derived assuming that it is possible to generate samples according to conditional laws when it is required. This is indeed a tough requirement but this problem is well identified and not particular to these randomised estimators (see Roberts 2011); especially Skilling (2006); Huber et al (2011); Guyader et al (2011) already acknowledge it and make use of Markov Chain Monte Carlo sampling. While a lot of ongoing work on nested sampling focus on improving these conditional simulations (*e.g.* Brewer et al 2011), in the present article we focus on theoretical statistical properties and suggest a possible solution to the issue of choosing a *bad* stopping criterion. Hence, it is out of the scope of the present work to benchmark nested sampling against other tailor-made methods such as Importance Sampling (see for example (Robert and Casella 2004) or (Glynn and Iglehart 1989)) on a list of specific cases.

The outline of this paper is as follows: Section 2 presents the common framework for rare event simulation and nested sampling and derives a new *ideal* (not practically implementable) estimator of $m = \mathrm{E}[X] = \mathrm{E}[g(\mathbf{U})]$. It is closely related to nested sampling with an infinite number of terms and is compared to the usual Monte Carlo estimator. Section 3 proposes two possible estimators based on the *ideal* one. Section 4 studies the specific case where $X = g(\mathbf{U})$ is heavy-tailed and Section 5 gives information on practical implementation and numerical results. Finally an Appendix gathers all the proofs.

2 Ideal estimator

From now on we consider a real-valued random variable X, which can be for instance the output of a mapping $X = g(\mathbf{U})$, as discussed in the Introduction.

Furthermore for a real-valued random variable X, one can write $X = X_+ - X_-$ with X_+ and X_- non-negative random variables. Then, $\mathrm{E}[X] = \mathrm{E}[X_+] - \mathrm{E}[X_-]$. Thus in the sequel and without loss of generality we assume that X is a non-negative random variable with law μ^X. We also assume that X has a continuous *cdf* F and we write p_x instead of $\mathrm{P}[X > x] = 1 - F(x)$, for any $x \in \mathbb{R}^+$.

2.1 Extreme event simulation

In this section we recast common results from (Huber et al 2011; Guyader et al 2011; Simonnet 2014) in a general framework.

Definition 1 (Increasing random walk) Let $X_0 = 0$ and define recursively the Markov sequence $(X_n)_n$ such that

$$\forall n \in \mathbb{N} : \mathrm{P}[X_{n+1} \in A \mid X_0, \cdots, X_n] = \frac{\mu^X(A \cap (X_n, +\infty))}{\mu^X((X_n, +\infty))}.$$

In other words $(X_n)_n$ is a strictly increasing sequence where each element is generated conditionally greater than the previous one. Considering the sequence $(T_n)_{n\geq 1}$ such that $T_n = -\log(\mathrm{P}[X > X_n])$, it can be shown that $(T_n)_{n\geq 1}$ is distributed as the arrival times of a Poisson Process with parameter 1. Thus, the counting random variable of the number of events before x: $M_x = \mathrm{card}\{n \geq 1 \mid X_n \leq x\}$ follows a Poisson law with parameter $t_x = -\log p_x$.

This result leads to the construction of a new estimator for the probability of exceeding a threshold x. Indeed Lehmann-Scheffé theorem states that the minimum-variance unbiased estimator (MVUE) for $p_x = e^{-t_x}$ is

$$\widehat{p_x} = \left(1 - \frac{1}{N}\right)^M \tag{1}$$

with $M = \sum_{i=1}^{N} M_x^i$ the sum of N *iid* realisations of M_x. Here we find back the LPA estimator, which means that LPA is only one possible practical implementation of this estimator; especially Walter (2015) shows that LPA generates a marked Poisson Process with parameter N. In any case, the statistical properties of $\widehat{p_x}$ are then well known:

Proposition 1 (Statistical properties of $\widehat{p_x}$)

$$\begin{aligned}
\mathrm{E}\left[\widehat{p_x}\right] &= p_x \\
\mathrm{var}\left[\widehat{p_x}\right] &= p_x^2\left(p_x^{-1/N} - 1\right)
\end{aligned}$$

This estimator exhibits a logarithmic efficiency and asymptotically achieves the Cramer-Rao bound $-p_x^2 \log p_x / N$. Comparing to classical Monte Carlo, it *replaces* the factor $1/p_x$ in the variance by $\log 1/p_x$ when $p_x \ll 1$ and $N \gg 1$:

	classical Monte Carlo	Poisson Process
Variance	$\frac{p_x(1-p_x)}{N}$	$p_x^2\left(p_x^{-1/N} - 1\right)$
Approx.	$\frac{p_x^2}{N}\frac{1}{p_x}$	$\frac{p_x^2}{N}\log\frac{1}{p_x}$

Remark 1 The MVUE of $t_x = -\log p_x$ is M/N. From this relation one could consider the suboptimal estimator for p_x:

$$\widetilde{p_x} = e^{-\frac{M}{N}} = \left(e^{-\frac{1}{N}}\right)^M. \tag{2}$$

From the moment-generating function of M we get the mean and variance of $\tilde{p}_x$:

$$\begin{aligned}
\mathrm{E}\left[\widetilde{p_x}\right] &= p_x^{N(1-e^{-1/N})} = p_x + \frac{-p_x \log p_x}{2N} + o\left(\frac{1}{N}\right) \\
\mathrm{var}\left[\widehat{p_x}\right] &= p_x^{N(1-e^{-2/N})} - p_x^{2N(1-e^{-1/N})} \\
&= \frac{-p_x^2 \log p_x}{N} + \frac{p_x^2 \log p_x}{N^2}\left(\log p_x + 1\right) + o\left(\frac{1}{N^2}\right).
\end{aligned}$$

Hence this suboptimal estimator has a positive bias of order $1/N$. The variances $\mathrm{var}\left[\widetilde{p_x}\right]$ and $\mathrm{var}\left[\widehat{p_x}\right]$ differ only from order $1/N^2$ and $\mathrm{var}\left[\widehat{p_x}\right] < \mathrm{var}\left[\widetilde{p_x}\right]$ as soon as $p_x < e^{-1}$.

2.2 Definition of the moment estimator

Noticing that for a non-negative real-valued random variable with mean $m = \mathrm{E}[X] = \mathrm{E}[g(\mathbf{U})]$ one has:

$$m = \int_0^\infty p_x \mathrm{d}x, \tag{3}$$

the idea is to use the optimal estimator of p_x (Eq. (1)) to build an estimator for m.

From now on we will assume that $N \geq 2$ point processes have been simulated and denote by $(M_x)_x$ the counting random variables associated with the marked Poisson Process: $\forall x > 0, M_x \sim \mathscr{P}(-N \log p_x)$. The sequence $(X_n)_{n\geq 1}$ is the cumulated one, *i.e.* the combination of the states of the N Markov Chains sorted in increasing order; then the associated $(T_n)_{n\geq 1}$ are the times of the marked Poisson Process with parameter N. We set $X_0 = 0$ and then consider the following estimator:

$$\begin{aligned}
\widehat{m} &= \int_0^\infty \left(1 - \frac{1}{N}\right)^{M_x} \mathrm{d}x \\
&= \sum_{i=0}^{\infty} \left(X_{i+1} - X_i\right)\left(1 - \frac{1}{N}\right)^i.
\end{aligned} \tag{4}$$

The second equality comes from the fact that $x \mapsto M_x$ is constant equal to i on each interval $[X_i, X_{i+1})$: there are 0 event before X_1, then 1 event before X_2, precisely at X_1, etc.

While the first form is easier to analyse because the law of $(M_x)_x$ is well determined, the second one paves the way for the practical implementation (see Section 3) and clarifies the link with Nested Sampling:

$$\widehat{m} = \sum_{i=1}^{\infty} X_i \left[\left(1 - \frac{1}{N}\right)^{i-1} - \left(1 - \frac{1}{N}\right)^i\right]. \tag{5}$$

This estimator is the limit of the nested sampling estimator with a deterministic scheme (Skilling 2006):

$$\widetilde{m} = \sum_{i=1}^{\infty} X_i \left(e^{\frac{1-i}{N}} - e^{\frac{-i}{N}}\right) \tag{6}$$

with slightly modified weights: $(1 - 1/N)$ instead of $e^{-1/N}$. This is a direct consequence of the fact that an optimal unbiased estimator for e^{-t_x} is not $e^{-\hat{t}_x}$ (see Section 2.1 Remark 1).

Proposition 2 (Statistical properties of $\widehat{m}$)

$$\mathrm{E}\left[\widehat{m}\right] = m \tag{7}$$

$$\mathrm{var}\left[\widehat{m}\right] = 2\int_0^\infty \int_0^x p_x p_{x'}^{1-1/N} \mathrm{d}x' \mathrm{d}x - m^2 \tag{8}$$

We thus have defined an unbiased estimator for m.

Remark 2 As a matter of comparison, $\widetilde{m}$ can also be written $\widetilde{m} = \int_0^\infty \widetilde{p_x} \mathrm{d}x$. Then Remark 1 allows us to conclude that $\tilde{m}$ has a positive bias of order 1/N.

Proposition 3 (Finiteness of $\mathrm{var}[\widehat{m}]$)

$$\forall N \geq 2,\ \mathrm{var}[\widehat{m}] \leq \frac{2}{1+1/N} E[X^{1+1/N}]^{2/(1+1/N)}.$$

Corollary 1 (Value of N) *Let $\varepsilon > 0$, if $\mathrm{E}\left[X^{1+\varepsilon}\right] < \infty$ then for any $N \geq 1/\varepsilon$, $\widehat{m}$ has a finite variance.*

While the usual Monte Carlo estimator requires the finiteness of $\mathrm{E}\left[X^2\right]$ to have a finite variance, this estimator only requires the finiteness of a moment of order $1+\varepsilon$. This is especially interesting when X is heavy-tailed and this case is further investigated in Section 4.

2.3 Comparison with classical Monte Carlo

As the finiteness condition of the variance of $\widehat{m}$ is much weaker than for a naive Monte Carlo estimator, one can expect a globally lower variance. This result is shown in Proposition 4. We first recall the crude Monte Carlo estimator:

$$\widehat{m}_{MC} \overset{\mathrm{def}}{=} \frac{1}{N}\sum_{i=1}^{N} X_i \tag{9}$$

with $(X_i)_i$ N *iid* random variables with law μ^X.

Proposition 4 *For any $N \geq 2$, $\mathrm{var}[\widehat{m}] \leq \mathrm{var}[\widehat{m}_{MC}]$.*

Thus the *ideal* nested sampling estimator (5) is always better than classical Monte Carlo in terms of variance and especially does not require the finiteness of the second-order moment of X to have a finite variance.

3 Randomised unbiased estimator

The ideal estimator (4) defined in Section 2 is not directly usable as it requires to simulate an infinite number of terms in sum (4). While the usual nested sampling implementations propose to stop the algorithm either after a given number of iterations, or according to some criterion estimated at each iteration, we propose a randomised unbiased estimator using recent results on paths simulation.

3.1 Definition

We are facing the issue of estimating $\mathrm{E}[\widehat{m}]$ while it is not possible to generate such a $\widehat{m}$ in a finite computer time. This problem is well identified in the field of Stochastic Differential Equations (SDE) where one often intends to compute the expectation of a path functional while only discrete-time approximations are available. Recently there have been two major breakthroughs that address this issue: first the Multilevel Monte Carlo (MLMC) method (Giles 2008) has introduced the idea of combining *intelligently* different biased estimators (levels of approximations) to speed up the convergence and reduce the bias; then McLeish (2011) and Rhee and Glynn (2013) have introduced a general approach to constructing unbiased estimator based on a family of biased ones. Basically in our context it randomises the number of simulated steps of the Markov chain, and slightly modifies the weights of the nested sampling to *remove the bias* of the final estimator.

More precisely let us consider the truncated estimators $(\widehat{m}_n)_{n\geq 1}$:

$$\widehat{m}_n = \int_0^{X_n} \left(1-\frac{1}{N}\right)^{M_x} \mathrm{d}x = \sum_{i=0}^{n-1} (X_{i+1}-X_i)\left(1-\frac{1}{N}\right)^i$$

and T a non-negative integer-valued random variable independent of $(X_n)_{n\in\mathbb{N}}$ such that $\forall i \in \mathbb{N}, \mathrm{P}[T \geq i] \overset{\mathrm{def}}{=} \beta_i > 0$; one builds the following estimator (with $\widehat{m}_0 = 0$):

$$\begin{aligned}\widehat{Z} &= \sum_{n=0}^{\infty} \frac{\widehat{m}_{n+1}-\widehat{m}_n}{\mathrm{P}[T\geq n]} \mathbb{1}_{T\geq n} = \sum_{n=0}^{T} \frac{\widehat{m}_{n+1}-\widehat{m}_n}{\mathrm{P}[T\geq n]} \\ &= \sum_{n=0}^{\infty} (X_{n+1}-X_n)\left(1-\frac{1}{N}\right)^n \frac{\mathbb{1}_{T\geq n}}{\mathrm{P}[T\geq n]}. \end{aligned} \tag{10}$$

Remark 3 The notation $\widehat{Z}$ might seem a bit confusing since Z is used in the Introduction for the evidence as in (Skilling 2006). This is to keep consistency with Rhee and Glynn (2013) notations where the randomising procedure comes from.

Proposition 5 (Statistical properties of $\widehat{Z}$)

$$\begin{aligned}\mathrm{E}\left[\widehat{Z}\right] &= m \\ \mathrm{var}\left[\widehat{Z}\right] &= \sum_{i=0}^{\infty} q_{i,N}\beta_i^{-1} - m^2\end{aligned}$$

with:

$$q_{i,N} = 2\left(1-\frac{1}{N}\right)^{2i} \int_0^{\infty}\int_{x'}^{\infty} p_x p_{x'}^{N-1} \frac{[-N\log p_{x'}]^i}{i!}\mathrm{d}x\mathrm{d}x'. \tag{11}$$

The asymptotic behaviour of the sequence $(q_{i,N})_i$ will drive the possible choices for the randomising distribution $(\beta_i)_i$: $\mathrm{var}\left[\widehat{Z}\right]$ to remain finite implies that $q_{i,N}\beta_i^{-1} \to 0$ when $i \to \infty$.

Lemma 1 *The sequence $(q_{i,N})_i$ goes to 0 at least at exponential rate. Furthermore, if X has density f_X such that $\|f_X\|_\infty < \infty$, it is also bounded from below by an exponentially decreasing sequence.*

Then it appears that the Geometric distribution plays a key role, as already stated by McLeish (2011). Hence we provide some theoretical results assuming that T is a geometric random variable.

Proposition 6 *If* $\mathrm{P}[T \geq n] = e^{-\beta n}$, $\beta > 0$, *then:*

$$\mathrm{var}\left[\widehat{Z}\right] = 2\int_0^\infty \int_0^x p_x p_{x'}^{1-\frac{1}{\gamma(\beta,N)}} \mathrm{d}x'\mathrm{d}x - m^2 \tag{12}$$

with $\gamma(\beta,N) = N/(1+(e^\beta-1)(N-1)^2)$.

This expression is indeed the same as the one of Proposition 2 with the function $\gamma(\beta,N)$ instead of N. Hence the greater γ the smaller $\mathrm{var}\left[\widehat{Z}\right]$. Furthermore one has directly all the results from Section 2.2, especially the finiteness conditions for the variance given in Proposition 3 and Corollary 1, replacing N by $\gamma(\beta,N)$.

While there is no value of β minimising $\mathrm{var}\left[\widehat{Z}\right]$ at a given N (the smaller β the smaller the variance of the randomised estimator $\widehat{Z}$), there is an optimal value for N at a given β, *i.e.* for a given finite computational budget: $N = \sqrt{1+\mathrm{E}[T]}$. One can reverse this relation, which gives:

$$\beta_{\mathrm{app}} \stackrel{\mathrm{def}}{=} \log\left(1+1/(N^2-1)\right). \tag{13}$$

Corollary 2 *Let* $N \geq 2$ *and* $\mathrm{P}[T \geq n] = e^{-n\beta_{app}(N)}$, *then:*

$$\mathrm{var}\left[\widehat{Z}\right](N) = \mathrm{var}\left[\widehat{m}\right]\left(\tfrac{N+1}{2}\right) \approx 2\,\mathrm{var}\left[\widehat{m}\right](N). \tag{14}$$

This means that instead of choosing an arbitrary stopping criterion for nested sampling, randomising the number of iterations and computing $\widehat{Z}$ allows for keeping an unbiased estimator without increasing drastically the variance (factor up to 2, reached with suboptimal implementation of Corollary 2). This result will be illustrated in the examples of Section 5.

3.2 Convergence rate

Throughout the paper we consider that the computational cost for generating a realisation of $\widehat{Z}$ is the number of simulated samples. Accordingly, in this section it is the number of calls to a simulator of a conditional law.

Proposition 7 *Let* τ *be the random variable of the number of samples required to generate* $\widehat{Z}$. *One has* $\tau = N+T$.

Corollary 3 (Convergence rate of $\widehat{Z}$) *For any non-negative integer-valued randomising variable* T *such that* $\mathrm{E}[T] < \infty$ *and* $\forall i \in \mathbb{N}$, $\mathrm{P}[T \geq i] > 0$, *one has:*

$$\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right] \geq 2q_{1,2} + O\left(\frac{1}{N}\right),\ N \to \infty. \tag{15}$$

If the inequality (15) is close to an equality then $\widehat{Z}$ has a canonical square-root convergence rate (as a function of the computational cost). However there is no guarantee on this rate of convergence. Especially Corollary 4 below shows that it is not the case when T has a geometric distribution.

Corollary 4 *If* T *is a Geometric random variable such that* $\forall n \in \mathbb{N}$, $\mathrm{P}[T \geq n] = e^{-\beta n}$ *with* $\beta = \Theta(1/N^{1+\varepsilon})$, $\varepsilon \geq 0$, *then:*

$$\begin{cases} \mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right] = \Theta(N) & \varepsilon \in [0,1] \\ \mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right] = \Theta(N^\varepsilon) & \varepsilon > 1. \end{cases}$$

Hence the unbiased randomised estimator of Corollary 2 with $\beta = \beta_{\mathrm{app}} = \Theta(1/N^2)$ does not have a canonical square-root convergence rate. Furthermore, even though the realisation of the geometric random variable gives a small number of iterations, one may want to run the algorithm longer to probe the tail of the likelihood function to make sure that no important part is missing (Skilling 2006). That is why the idea behind randomised estimators is to average several replicas of $\widehat{Z}$ because it will somehow average the quantities $\mathbb{1}_{T\geq n}/\mathrm{P}[T \geq n]$ in (10). More precisely, let $G(c)$ be the random variable of the number of simulations of $\widehat{Z}$ one can afford with a computational budget c:

$$G(c) = \max\{n \geq 0 \mid \sum_{i=1}^{n} \tau_i \leq c\}$$

where τ_i is the computational effort required to generate the i^{th}-sample $\widehat{Z}_i$, one considers the following estimator:

$$\widehat{\alpha}(c) = \frac{1}{G(c)} \sum_{i=1}^{G(c)} \widehat{Z}_i. \tag{16}$$

In this setting Glynn and Whitt (1992) showed a CLT-like result:

$$c^{1/2}(\widehat{\alpha}(c) - \mathrm{E}\left[\widehat{Z}\right]) \xrightarrow[c\to\infty]{\mathscr{L}} (\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right])^{1/2}\mathcal{N}(0,1). \tag{17}$$

Hence in our context one has to tune $(\beta_i)_i$ and N to minimise the product $\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right]$.

3.3 Optimal randomisation

Since T is a non-negative random variable one has $\mathrm{P}[T \geq 0] = \beta_0 = 1$. Let $\mathscr{C} = \{(\beta_i)_i \in (0,1]^{\mathbb{N}} \mid \beta_0 = 1 \text{ and } \forall i \in \mathbb{N}, \beta_{i+1} \leq \beta_i\}$; we intend to solve the optimisation problem:

$$\begin{aligned} &\underset{\substack{(\beta_i)_i\in\mathscr{C}\\ N\in[\![2,\infty)}}{\mathrm{argmin}}\ \mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right] \\ &\qquad = \underset{\substack{(\beta_i)_i\in\mathscr{C}\\ N\in[\![2,\infty)}}{\mathrm{argmin}} \left(N-1+\sum_{i=0}^{\infty}\beta_i\right)\left(\sum_{i=0}^{\infty} q_{i,N}\beta_i^{-1} - m^2\right) \end{aligned} \tag{18}$$

where the $(q_{i,N})_i$ are given by (11). Furthermore, one can rewrite the $(q_{i,N})_i$ assuming that X has a density $f_X > 0$. Indeed in this context X_n has a density f_n such that:

$$\forall n \geq 1,\ f_n(x) = N\frac{p_x^{N-1}(-N\log p_x)^{n-1}}{(n-1)!}f_X(x).$$

This gives:

$$\forall i \in \mathbb{N},\ q_{i,N} = \frac{2}{N}\left(1-\frac{1}{N}\right)^{2i}\mathrm{E}\left[\mathscr{R}(X_{i+1})\right]$$

with $\mathscr{R}(x) = \int_x^\infty p_u \mathrm{d}u / f_X(x)$. Hence we further assume that $(q_{i,N})_i$ is decreasing, which is the case for a Pareto random variable (see Section 4.1) and at least for any distribution for which $\mathscr{R}$ is non-increasing like exponential and uniform distributions. In this context Proposition 8 gives the optimal distribution for T for a given N.

Proposition 8 (Optimal distribution for T) *If $(q_{i,N})_{i\geq 1}$ is decreasing then the optimal distribution $(\beta_i^*)_i$ for T is given by:*

$$\forall i \in [\![0, i_0]\!], \beta_i^* = 1$$
$$\forall i > i_0, \beta_i^* = \sqrt{\frac{N+i_0}{S_0}}\sqrt{q_{i,N}}$$

with $i_0 = \min\{i \in \mathbb{N} \mid \sum_{j=0}^i q_{j,N} - m^2 > (N+i)q_{(i+1),N}\}$ and $S_0 = \sum_{j=0}^{i_0} q_{j,N} - m^2$.

It is part of the proof in the appendix that i_0 is well defined and so it appears that the optimal distribution enforces the estimator to go at least until the i_0^{th} event. Recalling $(X_n)_n$ is the cumulated Markov Chain (associated with the marked Poisson Process with parameter N), this can be understood in the sense that on average, at least N events are necessary to *use* at least one time each process. Even if the link between i_0 and N is not that straightforward, one can then conjecture that $\lim_{N\to\infty} i_0 = \infty$.

Corollary 5 (Bounds on β_i^*) *For all $i > i_0$, one has:*

$$\sqrt{\frac{q_{i,N}}{q_{i_0+1,N}}} > \beta_i^* \geq \sqrt{\frac{q_{i,N}}{q_{i_0,N}}}. \tag{19}$$

Thus the tail of the optimal distribution $(\beta_i^*)_i$ is exponentially decreasing by Lemma 1. From these bounds on the $(\beta_i)_i$ one can also derive bounds on the variance:

$$q_{i_0+1,N}\,\mathrm{E}[\tau]^2 < \mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right] \leq q_{i_0,N}\,\mathrm{E}[\tau]^2.$$

Assuming $\lim_{N\to\infty} i_0 = \infty$ and using the lower bound on $q_{i,N}$ from Lemma 1, one can show that $\lim_{N\to\infty}\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right] = \infty$, which implies the existence of an optimal N. Section 4.1 presents an exact resolution of this optimisation problem for a Pareto random variable.

Finally, we have presented in this section the framework for an optimal resolution of Problem (18) and proven existence of a solution under reasonable assumptions ($(q_{i,N})_i$ is decreasing and $\lim_{N\to\infty} i_0 = \infty$). Furthermore the comprehensive resolution in the case of a Pareto distribution in Section 4.1 legitimises these assumptions. Generally speaking, if $(q_{i,N})_{i\geq 1}$ is not decreasing the optimisation has to be performed over all the decreasing sub-sequences of $(q_{i,N})_i$, which turns it into a combinatorial problem (see Rhee and Glynn 2013, Theorem 3).

3.4 Geometric randomisation

On the one hand the computation of the optimal distribution for T can be quite demanding in computer time; and on the other hand the geometric law plays a key role as for any distribution p_x, the sequence $(q_{i,N})_i$ decreases at exponential rate and the optimal randomising distribution (when $(q_{i,N})_i$ is decreasing) is somehow a shifted geometric law. Therefore we study the parametric case where $\mathrm{P}[T \geq n] = e^{-\beta n}$, $\beta > 0$ and tune β and N to minimise $\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right]$.

Using the exponential power series in $\mathrm{var}\left[\widehat{Z}\right]$ (cf. Eq. (12)), the optimisation problem (18) becomes:

$$\min_{\substack{\beta>0\\ N\in[\![2,\infty)}} \left(N + \frac{1}{e^\beta - 1}\right)\left(\sum_{i=0}^\infty q_{i,2}\left(\frac{2}{\gamma(\beta,N)}\right)^i - m^2\right). \tag{20}$$

Proposition 9 *There exists a global minimiser (β_{opt}, N_{opt}) to Problem* (20)*. Furthermore, (β_{opt}, N_{opt}) satisfies the relationship:*

$$\beta_{opt} = \log\left(1 + \frac{2}{N_{opt}^2 - 1 + (N_{opt}-1)\sqrt{N_{opt}^2 + 6N_{opt} + 1}}\right). \tag{21}$$

Hence there is always an optimal solution to Problem (20), meaning this parametrisation is meaningful.

To summarise we have shown that by randomising the finite number of iterations and slightly modifying the weights of the original nested sampling, it is possible to define an unbiased estimator for the mean of any real-valued random variable with continuous *cdf*, resolving the issue of choosing an *appropriate* stopping criterion. With a suboptimal geometric randomisation as in Corollary 2, the variance is at most twice the one of the ideal case (estimator (4)). However it is not usable with a fixed predetermined computational budget and its convergence rate is slower than the canonical square-root one. To circumvent this limitation, the idea is to average several replicas of the randomised unbiased estimator (see Eq. (16)). This new estimator remains unbiased and also supports a Central Limit Theorem.

All these theoretical results assume that it is possible to generate conditional random variables when required, as for the original nested sampling algorithm (see Skilling 2006, Section 9). Efficient conditional simulation can be carried out in different ways, from perfect simulation (see for example Propp and Wilson 1996) to approximation using random walk Metropolis-Hastings. The aim of this paper is not to challenge this hypothesis in a general manner but only to provide a new insight on the risk of choosing a *bad* stopping criterion in nested sampling, and to propose an other tool to deal with this issue. Since nested sampling has been applied successfully to a great number of problems so far, these results are expected to hold in these situations. Also the examples of Section 5 are in good agreement with these theoretical results.

In the next section, we discuss the different stopping criteria usually recommended for nested sampling and parallel implementation of the estimators.

3.5 Parallel implementation

Skilling (2006, Section 7) presents two possible termination rules based on criteria evaluated on-the-fly:

- stop when the greatest expected increment (current weight and biggest found likelihood value) is smaller than a given fraction of the current estimate;
- stop when the number of iterations *significantly* exceeds NH with H the information, estimated on-the-fly.

Chopin and Robert (2010) use an other stopping criterion, close to the first one above, it is: "stop when the new increment is smaller than a given fraction of the current estimate". An other option is to do a predetermined number of iterations (Brewer et al 2011). Unfortunately these criteria give no guarantee on the convergence of the estimator to the sought value and may lead to biased estimation.

A first difference between the three first criteria and the last one stands in the fact that this latter uses a known computational budget while the others ones will run until the criterion is satisfied; hence there is no way to estimate the (random) final number of iteration in advance. This difference is also to be found between $\widehat{Z}$ and $\widehat{\alpha}$: the first one will use a random number of simulated samples (the draw of the randomising variable) while the second one is defined with a fixed computational budget. Hence these two categories of estimators cannot be compared because the setting is not the same.

An other main difference between these estimators is whether they enable parallel computation or not. The three first stopping criteria need to be evaluated at each iteration and are based on quantities estimated with the full process with parameter N. Hence they do not allow for parallel computation. On the other hand, with a predetermined total number of iterations, parallel computation on the model of (Walter 2015, Section 4.2) can be carried out. The randomised estimator $\widehat{Z}$ also enables this feature as the random number of iterations is drawn before the algorithm starts. Considering $\widehat{\alpha}$, each replica can be computed in parallel, and further the computation of each replica also allows for parallel implementation. Hence $\widehat{\alpha}$ allows for a *double* parallelisation, which is worth noticing as it may require a substantial computational budget to become effectively Gaussian.

To conclude, one stresses out the fact that among estimators with random computational budget, $\widehat{Z}$ is the only one allowing for parallel computation; furthermore it is also the only one unbiased and its variance is at worst twice the one of the ideal estimator (upper bound reached with suboptimal implementation of $\widehat{Z}$ as in Corollary 2). Both fixed-budget estimators enable parallel implementation; however nested sampling with a predetermined number of iterations has no reason to be close to the sought value. On the other hand, $\widehat{\alpha}$ is unbiased and supports a CLT. All these considerations are illustrated in Section 5.

4 Application to heavy-tailed random variables

In this section we give insights on the properties of the new estimator when $X = g(\mathbf{U})$ is heavy-tailed. Mean estimation for heavy-tailed random variables is a well identified problem often addressed by some parametric assumptions on the *cdf* of X; see Beirlant et al (2012) for a comprehensive overview of tail index estimation, and Peng (2001); Johansson (2003); Necir et al (2010); Hill (2013) for references on mean estimation for heavy-tailed random variables.

In the sequel we then give explicit results for the Pareto distribution $p_x = \mathrm{P}[X > x] = 1 \wedge x^{-a}$, $a > 1$.

4.1 Exact resolution for a Pareto distribution

With an analytic form for the *cdf* of X, we can derive explicit formulae for the variance (Eq. (8)) and the optimisation problem (18).

First we compare the variance of the ideal estimator $\widehat{m}$ against usual Monte Carlo and Importance Sampling estimators. In this latter case the importance density is chosen to be a Pareto distribution with parameter $b > 0$.

Proposition 10 (Variance comparison) *For a Pareto distribution, one has $m = a/(a-1)$ and the variances write:*

$$a > 2,\ \mathrm{var}[\widehat{m}_{MC}] = \frac{m(m-1)^2}{2N - mN}$$

$$a > \frac{2N}{2N-1},\ \mathrm{var}[\widehat{m}] = \frac{m(m-1)^2}{2N - m}$$

$$a > 1 + \frac{b}{2},\ \mathrm{var}[\widehat{m}_{IS}] = \frac{m^2(B-1)^2}{N(2B-1)}$$

with $B = (a-1)/b \in (1/2, \infty)$.

It is clearly visible that the classical Monte Carlo estimator needs a second-order moment while $\widehat{m}$ only requires $a > 2N/(2N-1) \approx 1 + 1/2N$ and $\widehat{m}_{IS}$ requires $a > 1 + b/2$; it also illustrates the result of Proposition 4: $\mathrm{var}[\widehat{m}] < \mathrm{var}[\widehat{m}_{MC}]$. The optimal value $b = a-1$ cancels out $\mathrm{var}[\widehat{m}_{IS}]$. It is well known that there is an optimal density q for IS that cancels out the variance of the IS estimator but it is case-specific: here a Pareto density with parameter $a-1$.

Remark 4 (Limit distribution of classical Monte Carlo estimator) In the case of Pareto distribution, when $a > 2$ the Central Limit Theorem gives the limit law of the estimator while for $1 < a < 2$ the Generalised Central Limit Theorem (see for example Embrechts et al 1997) states that $\sum_i X_i$ is in the domain of attraction of a stable law with parameter a:

$$N^{1-1/a}\left(\frac{1}{N}\sum_{i=1}^{N} X_i - m\right)\frac{1}{C_a} \xrightarrow[N\to\infty]{\mathscr{L}} X_a$$

with the characteristic function of X_a, ϕ_{X_a}, writing $\phi_{X_a}(t) = \exp\left[-|t|^a\left(1 - i\left(\tan\left(\pi a/2\right)\right) sgn(t)\right)\right]$ and C_a the normalising constant $C_a = \pi^{1/a}\left(2\Gamma(a)\sin \pi a/2\right)^{-1/a}$.

We now detail the resolution of optimisation problems (18) and (20). Especially we first explicit the form of the sequence $(q_{i,N})_i$ defined in Eq. (11).

Proposition 11 *If X is a Pareto random variable with parameter $a > 1$, then:*

$$\forall i \in \mathbb{N}, q_{i,N} = \frac{2}{(a-1)(aN-2)}\left[\frac{a(N-1)^2}{N(aN-2)}\right]^i + \frac{\mathbb{1}_{i=0}(a+1)}{2(a-1)}.$$

Hence for a Pareto distribution $(q_{i,N})_i$ is decreasing. One can then look for i_0, the solution of the problem $i_0 = \min\{i \in \mathbb{N} \mid \sum_{j=0}^{i} q_{j,N} - m^2 > (N+i)q_{(i+1),N}\}$. Whilst an exact solution can be expressed using the lower branch of the Lambert W function (see for example Corless et al 1996), the following proposition gives an asymptotic approximation when $N \to \infty$ to precise the growth rate of i_0.

Proposition 12 *If X is a Pareto random variable, then:*

$$i_0 = \frac{Nm}{2}\left(\log N + \log\log N - \log(\frac{m}{2})\right) + o(N),\ N \to \infty.$$

Corollary 6 (Order of magnitude of $\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right]$)

$$\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right] \underset{N\to\infty}{\sim} \left(\frac{m(m-1)}{2}\right)^2 \log N.$$

Corollary 6 shows that $\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right] \to \infty$ when $N \to \infty$ so there is an optimal value for N that minimises $\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right]$; a numerical resolution for several values of a from 1 to 3 was performed and the result is displayed in Figure 1a. We also present in Figure 1b a comparison between the optimal variance (with the optimal distribution $(\beta_i^*)_i$ and optimal N) and the classical Monte Carlo one. There we can see that for $a \lesssim 2.5$ the new estimator (16) performs better in terms of variance; especially for $a < 2$ it remains finite while $\mathrm{var}[\widehat{m}_{MC}] = \infty$.

As explained in Section 3.4 we consider now a Geometric random variable T with parameter β for the random truncation.

Proposition 13 *If X is a Pareto random variable with parameter $a > 1$ and $\forall n \in \mathbb{N}$, $\mathrm{P}[T \geq n] = e^{-\beta n}$ then:*

$$\mathrm{var}\left[\widehat{Z}\right] = \frac{m(m-1)^2}{2\gamma(\beta, N) - m}$$

and

$$\beta_{opt} = \log\left(\frac{1}{B_+} + 1\right) \tag{22}$$

where B_+ is the positive root of the quadratic polynomial $P(B)$:

$$P(B) = \frac{2N_{opt} - m}{(N_{opt}-1)^2}B^2 - 2mB - \left(m(N_{opt}-1)^2 + 2N_{opt}^2\right).$$

With this relation and the one of Eq. (21) one can derive the optimal parameters $(\beta_{\mathrm{opt}}, N_{\mathrm{opt}})$. Figure 1a shows a numerical resolution of this problem for several values of $a \in (1, 3]$.

Furthermore, if one considers the approximation of the optimisation problem (20) with relation (13) instead of (21), one has to minimise $N \mapsto (N^2+N-1)m(m-1)^2/(N+1-m)$. Denoting N_{app} this minimiser, one has:

$$N_{\mathrm{app}} = \max\left(m - 1 + \sqrt{m^2 - m - 1}, 2\right) \tag{23}$$

This approximation is the red dotted-dashed line of Figure 1a. As we can see, it is in good agreement with the optimal values, both for the parameter N and for the global variance (see further Section 4.2 and Figure 1b).

4.2 Comparison of the estimators

We have seen in Sections 3.3 and 3.4 two ways of implementing the *ideal* estimator $\widehat{m}$ defined in Section 2.2 with a fixed computational budget. Then we have presented their exact behaviours in a case of a Pareto random variable. These two ways involve a truncation of the infinite sum (4) by an integer-valued random variable T. In the first implementation the distribution of T and the number N of point processes are optimised in order to minimise the estimator variance. In the second implementation, the distribution of T is enforced to be geometric and its parameter as well as N are optimised.

While the first implementation is optimal in terms of variance, it requires to solve a combinatorial problem, which can turn it into a poorer algorithm in terms of computational time. In this scope, the parametric algorithm constraining the randomising variable T to be geometric with parameter β is much simpler to implement. The aim of this section is to benchmark these two implementations and to challenge the optimal parameters against the fixed ones we will suggest.

More precisely, while both optimisations ended up with optimal parameters depending on the distribution of X, we also consider the parametric algorithm with parameter β_{app} given by (13) and $N = N_{\text{app}}$, 2, 5 or 10.

Figure 1b shows the relative increase of the standard deviations due to the suboptimal implementations for a given computational budget, *i.e.* for a given number of generated samples. It also shows the standard deviation ratios between the optimal implementation, the classical Monte Carlo estimator (9) and $\widehat{m}$ given by (4). For this latter, it is assumed that its computational cost is N, *i.e.* that it costs 1 to simulate an increasing random walk (see Definition 1) while it requires an infinite number of simulated samples. This calls for certain comments:

- the parametric implementation with optimised parameters $(\beta_{\text{opt}}, N_{\text{opt}})$ remains competitive against the optimal implementation (solid black line going from ≈ 1.3 to ≈ 1.1);
- the parametric implementation with parameters β_{app} and N_{app} is almost not distinguishable from the parametric implementation with optimal parameters β_{opt} and N_{opt}. This means that it is not necessary to strive to estimate the parameters $(\beta_{\text{opt}}, N_{\text{opt}})$;
- the classical Monte Carlo estimator is better than the optimal implementation as soon as $a \gtrsim 2.5$ and better than the parametric implementation as soon as $a \gtrsim 2.3$; this confirms that nested sampling is especially convenient for heavy-tailed random variables;
- the standard deviation of $\widehat{m}$ illustrates the efficiency of the ideal estimator compared to the classical Monte Carlo one (cf. Proposition 10), with a standard deviation at least twice as small;
- generally speaking and without any knowledge on the distribution of X, N should not be set too small as the variance increases much faster when it is smaller than the optimal value; especially with $\beta = \beta_{\text{app}}$ finiteness condition of the variance writes $a > 1 + 1/N$.

Given these results we can consider that the parametric implementation is a good trade-off between minimal variance estimation and complexity, especially when no information on the distribution of X is provided.

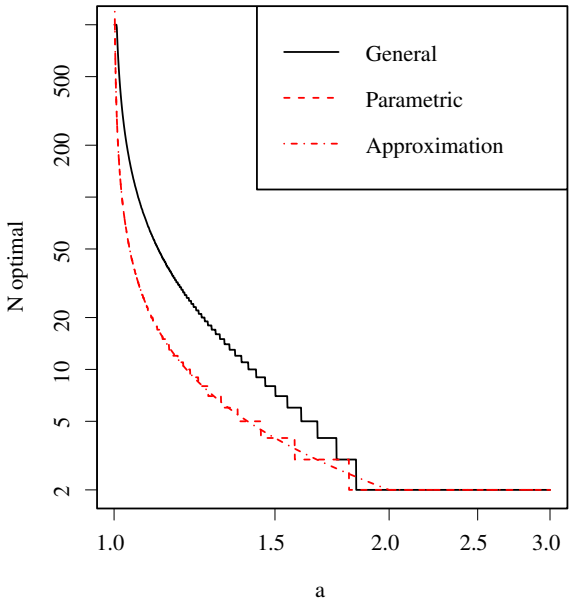

(a) Optimal values for N in the general (cf Section 3.3) and in the parametric (cf Section 3.4) cases with the approximation of equation (23).a is the parameter of the Pareto distribution.

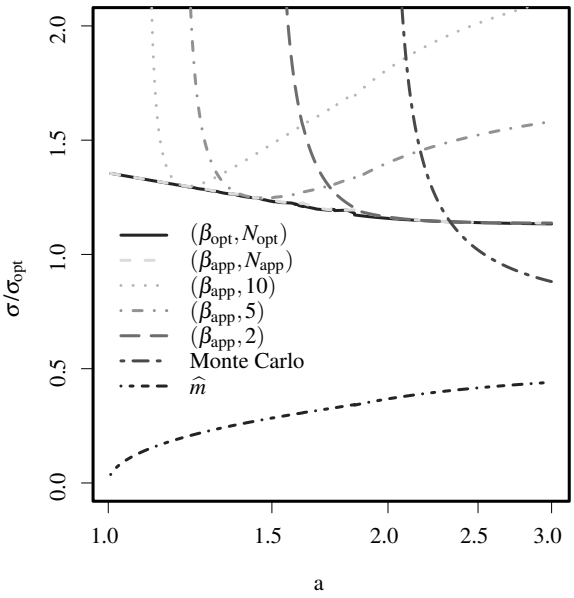

(b) Ratios of the standard deviations of different estimators over the standard deviation of the optimal estimator $\widehat{\alpha}$ of Section 3.3. The classical Monte Carlo estimator is defined in Eq. (9); $\widehat{m}$ is the ideal estimator (4); the other estimators are randomised estimators (16) with enforced geometric distribution for T with parameter β and N as follows: $(\beta_{\text{opt}}, N_{\text{opt}})$: optimal parameters of Proposition 9; $(\beta_{\text{app}}, N_{\text{app}})$: approximated optimal parameters of Eq. (13) and (23). a is the parameter of the Pareto distribution.

Fig. 1: Theoretical resolution of problems (18) and (20) when $\mathrm{P}[X > x] = 1 \wedge x^{-a}$.

5 Example

The aim of this section is to check the consistency between theoretical formulae and practical results with non-ideal con-

ditional sampling. It is also to demonstrate how *bad* stopping criteria can alter nested sampling and how randomised estimators can resolve this issue. We first explain how we perform conditional simulation and give pseudo-code for both $\widehat{Z}$ and $\widehat{\alpha}$. Then we present results on an example from (Skilling 2006, Section 18) that we slightly modify. The presented results are obtained with 500 simulations and boxplots extend to the extreme values.

5.1 Simulating conditional distributions

When no conditional sampler is available, a general idea is to use convergence properties of an ergodic Markov Chain to its unique invariant probability distribution. Assuming $\mathbf{U}$ is a d-dimensional random vector with *pdf* f_U, it means that we intend to generate a Markov Chain with stationary *pdf* $\propto \mathbb{1}_{g(\mathbf{u})>x} f_U(\mathbf{u})$. This implementation is rather simple when a reversible transition kernel is available. In the sequel we make use of the transition kernel suggested by Cérou et al (2012) detailed on Algorithm 1 for Gaussian input space.

Algorithm 1 Transition kernel for $\mathbf{U} \overset{\mathscr{L}}{\sim} \mathscr{N}(0, \mathbf{I}_d)$ (Cérou et al 2012; Guyader et al 2011)

Require: initial state $\mathbf{u}$, σ, *burn-in* b
  while $b > 0$ **do**
    Pick $\mathbf{W}$ from a standard multivariate Gaussian distribution
    $\mathbf{U}^* \leftarrow \dfrac{\mathbf{u} + \sigma\mathbf{W}}{\sqrt{1+\sigma^2}}$
    if $g(\mathbf{U}^*) > x$ **then**
      $\mathbf{u} \leftarrow \mathbf{U}^*$
    end if
    $b \leftarrow b - 1$
  end while
  return u

Because the goal is to reach the stationary state of the Markov Chain, several transitions have to be done to insure independence between the starting point and the final sample and adequacy with the targeted distribution. This number of transitions is referred to as a *burn-in* parameter b. Eventually the last generated sample is kept. In theory, one can start from any point provided the *burn-in* is large enough but practically speaking it is profitable to start with a point approximately following the targeted distribution as *burn-in* will then serve mainly independence purpose. Furthermore, the step size σ is initialised at $\sigma = 0.3$ and further updated after each use of the transition kernel – *i.e.* each b transitions – to get an acceptance rate close to 0.5.

Remark 5 The *burn-in* parameter increases the cost of an estimator because it needs several simulations for only one sample. In this context, the computational cost defined in Proposition 7 becomes $\tau = N + bT$ and is the number of calls to the generator of X (which amounts to generate $\mathbf{U}$ and to call g). Since this increase is common to all algorithms considered here, we will not mention it any more.

5.2 Pseudo-code

As explained above, we do not intend to solve the combinatorial optimisation problem in the general case and so we present here a pseudo-code for the parametric case. Reader interested in the optimal resolution is referred to (Rhee and Glynn 2013). We then present in Algorithm 2 how to compute $\widehat{Z}$ and in Algorithm 3 how to compute $\widehat{\alpha}(c)$. In this latter case we assume that N and β are given, being optimised (with previous knowledge or simulations) or not.

Algorithm 2 Pseudo-code for $\widehat{Z}$

Require: N, β
  Generate $\mathtt{T}$ according to $\mathrm{P}[T \geq n] = e^{-\beta n}$
  Generate N random variables $(X_i)_{i=1..N}$ according to μ^X
  $\mathtt{times}[0] \leftarrow 0$; $\mathtt{delta}[0] \leftarrow 0$
  for i in 1:T **do**
    $\mathtt{ind} \leftarrow \mathrm{argmin}_j X_j$
    $\mathtt{times}[\mathtt{i}] \leftarrow X_{\mathtt{ind}}$
    $\mathtt{delta}[\mathtt{i}-1] \leftarrow (\mathtt{times}[\mathtt{i}] - \mathtt{times}[\mathtt{i}-1]) \cdot \dfrac{(1-1/N)^i}{e^{-\beta i}}$
    Generate $X^* \sim \mu^X(\cdot \mid X > X_{\mathtt{ind}})$
    $X_{\mathtt{ind}} \leftarrow X^*$
  end for
  $\mathtt{ind} \leftarrow \mathrm{argmin}_i X_i$
  $\mathtt{times}[\mathtt{T}+1] \leftarrow X_{\mathtt{ind}}$
  $\mathtt{delta}[\mathtt{T}] \leftarrow (\mathtt{times}[\mathtt{T}+1] - \mathtt{times}[\mathtt{T}]) \cdot \dfrac{(1-1/N)^T}{e^{-\beta T}}$
  $\widehat{Z} = \sum_{i=0}^{T} \mathtt{delta}[\mathtt{i}]$

Remark 6 Note that in Algorithm 2, N is both the theoretical parameter of the number of increasing random walks per $\widehat{Z}$ and the size of the population for conditional simulation purpose. Hence it should not be set too small according to the dimension of the problem. This is a side effect of this practical implementation. Alternatively one could generate several $\widehat{Z}_i$ sequentially to *aggregate* all the samples for conditional simulations. Hence N could be chosen only according to theoretical guidelines. However it would disable parallel implementation. Some recent work on the parallel implementation of Sequential Monte Carlo may be used here (Vergé et al 2013). Note also that it is not necessary to consider only the minimum of the N samples in Algorithm 2; however in the context of Markov Chain drawing it is better to select the starting point in a relatively big population already following the targeted distribution.

Basically, Algorithm 3 is just a wrap-up of Algorithm 2 with an update of the remaining computational budget. If

Algorithm 3 Pseudo-code for $\widehat{\alpha}(c)$

Require: c, N, β
  $G \leftarrow 0$; $\widehat{\alpha} \leftarrow 0$;
  while $c > 0$ **do**
    Generate T^* according to $\mathrm{P}[T \geq n] = e^{-\beta n}$
    $c = c - (N + T^*)$; $G = G + 1$; $T[G] = T^*$
  end while
  if $c < 0$ **then** ▷ discard the last replica if it exceeds the budget
    $G = G - 1$; $T = T[1:G]$
  end if
  for g in 1:G **do**
    Start Algorithm 2 from step 2 with $\mathtt{T} = T[g]$
    $\widehat{\alpha} = \widehat{\alpha} + \widehat{Z}$
  end for
  $\widehat{\alpha} = \widehat{\alpha}/G$

one intends to use Markov Chain simulation as presented in Section 5.1 then one has to take into account the *burn-in* b and update c in Algorithm 3 as follows: $c = c - (N + bT^*)$.

5.3 Variance increase

In this section, we intend to check the variance increase between the ideal estimator $\widehat{m}$ of Section 2.2 and the suboptimal randomised estimator of Corollary 2. To do so, we use an example from Skilling (2006) where it is known that 100 iterations per particle on average are enough. We also compute (NS) the original nested sampling estimator, *i.e.* the estimator of Eq. 6. (NS) and $\widehat{m}$ differ only in the weights used: $\exp -1/N$ instead of $1 - 1/N$; thus they are computed in the same run. The aim is to estimate the evidence of a likelihood with uniform prior over a $d-$dimensional unit cube: $m = \mathrm{E}[g(\mathbf{U})] = \mathrm{E}[X]$ with:

$$g(\mathbf{u}) = 100 \prod_{i=1}^{d} \frac{e^{-u_i^2/2u^2}}{\sqrt{2\pi}u} + \prod_{i=1}^{d} \frac{e^{-u_i^2/2v^2}}{\sqrt{2\pi}v}, \tag{24}$$

$\mathbf{U} \sim \mathscr{U}\left(-[\frac{1}{2}, \frac{1}{2}]^d\right)$, $d = 20$, $u = 0.01$ and $v = 0.1$. This represents a Gaussian "spike" of width 0.01 superposed on a Gaussian "plateau" of width 0.1. Figure 2 plots the log-likelihood $\log x$ against the log-tail distribution $\log p_x$.

We then run nested sampling with stopping criterion "number of iterations = $100N$" as well as $\widehat{Z}$ for several values of N from 100 to 500. Figure 3 shows the boxplots of the estimators. On the one hand $\widehat{Z}$ has good convergence properties, on the other hand the bias and variance increase due to the original nested sampling weights is clearly visible. Table 1 summarises these numerical results: both $\widehat{Z}$ and $\widehat{m}$ are unbiased while (NS) has a bias of order $1/N$ (cf Remark 2). The variance increase between $\widehat{m}$ and $\widehat{Z}$ is in good agreement with the theoretical relationship of Corollary 2, it is $\mathrm{var}\left[\widehat{Z}\right](N) = \mathrm{var}[\widehat{m}]\,((N+1)/2) \approx 2\,\mathrm{var}[\widehat{m}]$. Also the ratio $\mathrm{var}[\mathrm{NS}]\,/\,\mathrm{var}[\widehat{m}]$ goes from 1.14 to 1.9. This variance increase appears to be of order $1/N^2$, which is consistent with the variance increase between $\widehat{p_x}$ and $\widetilde{p_x}$ (see Remark 1).

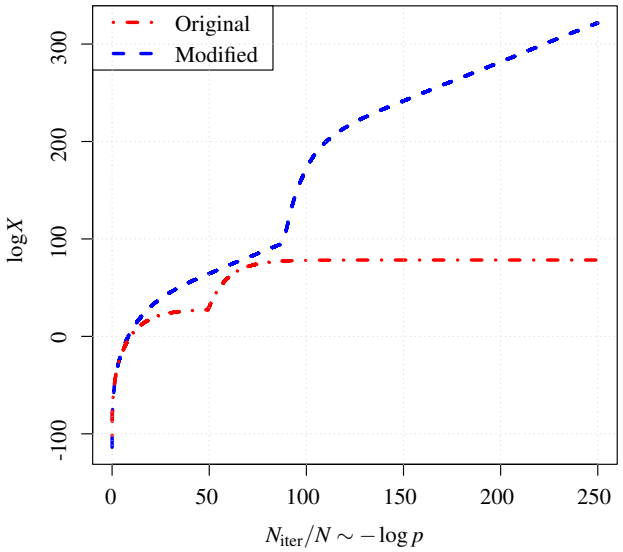

Fig. 2: Log-Likelihood against probability for the original example of Skilling (2006, Section 18) (Eq. (24)) and the modified version (Eq. (25)). Both lines are got from a sample run of nested sampling with $N = 300$ and stopping criterion $250N$ iterations.

Hence, the optimal choice of the nested sampling weights leads to significant variance reduction and removes the bias of the original nested sampling when it goes *far enough*. Unbiasedness can be maintained at the cost of at most doubling the variance of the estimator and even less compared to the currently used nested sampling weights. Furthermore, there is no need to choose (and justify) a stopping criterion for nested sampling any more.

N	100	200	300	400	500
$\mathrm{E}[\mathrm{NS}]$	142.3	117.7	114.6	111.5	109.5
$\mathrm{E}[\widehat{m}]$	103.0	100.8	102.8	102.8	102.6
$\mathrm{E}\left[\widehat{Z}\right]$	111.9	97.4	100.4	103.7	102.4
$\mathrm{var}\left[\widehat{Z}\right]/\mathrm{var}[\widehat{m}]$	3.23	2.49	1.90	2.20	1.70
$\mathrm{var}[\mathrm{NS}]/\mathrm{var}[\widehat{m}]$	1.90	1.33	1.24	1.17	1.14
$\mathrm{var}\left[\widehat{Z}\right]/\mathrm{var}[\mathrm{NS}]$	1.71	1.87	1.54	1.87	1.5

Table 1: Variance increase between the randomised unbiased nested sampling estimator $\widehat{Z}$, the original biased nested sampling (NS) and the ideal unbiased estimator $\widehat{m}$.

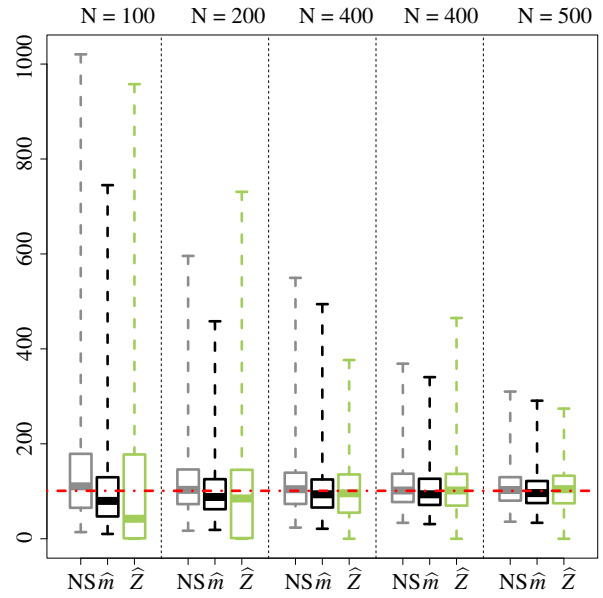

Fig. 3: Boxplots of ideal *infinite* nested sampling $\widehat{m}$ of Eq. (4) and (NS) of Eq. (5) and randomly truncated $\widehat{Z}$ (Corollary 2) for the estimation of $\mathrm{E}[g(\mathbf{U})]$ with g as in Eq. (24) and $\mathbf{U} \sim \mathscr{U}\left(-[\frac{1}{2},\frac{1}{2}]^d\right)$, $d = 20$. (NS) Ideal nested sampling is got with $N_{\text{iter}} = 100N$ as this is known to be enough in this case. (NS) and $\widehat{m}$ are obtained from the same runs. The (red) dot-dashed line is the theoretical value of m.

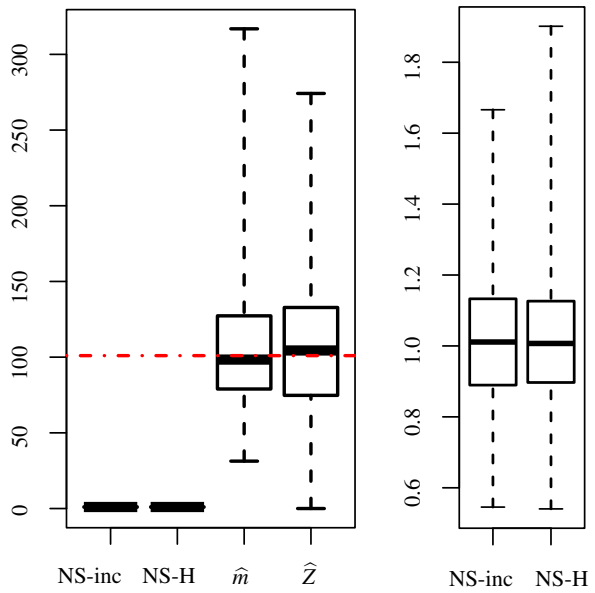

(a) Nested sampling estimators with adaptive stopping criteria, $\widehat{m}$ and $\widehat{Z}$

(b) Zoom on the nested sampling estimators with adaptive stopping criteria

Fig. 4: Effect of the choice of a stopping criterion for nested sampling estimator when estimating $\mathrm{E}[g(\mathbf{U})]$ with g as in Eq. (24) and $\mathbf{U} \sim \mathscr{U}\left(-[\frac{1}{2},\frac{1}{2}]^d\right)$, $d = 20$. (NS-inc): nested sampling stopped when current increment is less than 10^{-8} times the current estimator; (NS-H): nested sampling stopped when the number of iterations exceeds $2NH$; $\widehat{m}$ and $\widehat{Z}$ as in Figure 3. The (red) dot-dashed line is the theoretical value of m.

5.4 Adaptive stopping criteria

As we stated in the Introduction, one of the main concern of this paper was to point out the potential risk of using nested sampling with a *bad* stopping criterion. In this context we run nested sampling on the previous example with the adaptive stopping criteria mentioned in Section 3.5. The first one is directly picked out from (Chopin and Robert 2010), it is "stop when the current increment is less than 10^{-8} times the current estimate". The second one is based on the estimation of the information H and is the one described in the Appendix of (Skilling 2006); it is "stop when the number of iterations is greater than $2NH$". Figure 3 shows that for $N = 500$ the estimators should be well converged and so we set $N = 500$.

Figure 4 shows that nested sampling estimator can be not consistent if the termination rule is not well-chosen. Here both implementations miss the spike. In this context, the random truncation of $\widehat{Z}$ appears as a conservative practice. However, even though $\widehat{Z}$ allows for parallel computing (cf. Walter 2015, Section 4.2), $\widehat{Z}$ as well as the adaptive stopping criteria do not let work with a fixed computational budget. Yet one may have to work with fixed computational resources.

5.5 Nested sampling with fixed computational budget

There is only one nested sampling implementation which allows for fixing the total computational budget in advance. It is the one which stops after a given number of iterations. Following Rhee and Glynn (2013) we have proposed in Sections 3.3 and 3.4 a randomised estimator which also works with a predetermined computational budget. It is still unbiased and supports a Central Limit Theorem. The goal of this section is to compare these two estimators. We slightly modify the previous example (24) to narrow the spike: $u = 0.001$ instead of $u = 0.01$, and to make the random variable heavy-tailed:

$$g_{\text{ht}}(\mathbf{u}) = g(\mathbf{u}) \Big/ \left(\sum_{i=1}^{d} u_i^2\right)^{0.4d}. \tag{25}$$

Figure 2 compares this modified example with the original one. The heavy-tailed behaviour with tail index $1/0.8 = 1.25$ is clearly visible (limit slope of log-likelihood is 0.8) as well as the effect of the narrower spike (shift of the mass from $-\log p \approx 50$ to $-\log p \approx 90$). With Inv-χ^2 approximation of $1/\sum U_i^2$, the sought value is $\mathrm{E}[g_{\text{ht}}(\mathbf{U})] \approx 1.08 \times 10^{42}$.

Nested sampling is run with $N = 1000$ and $N = 10000$. We stop it after $100N$ iterations as in (Brewer et al 2011). This makes a total computational budget $c = 10^5$ (resp. 10^6). $\widehat{\alpha}$ is implemented with a suboptimal geometric randomising variable with parameter β_{app} (Eq. (13)) and $N = 20$. According to Remark 6, $N = d$ because it is both the theoretical parameter of $\widehat{\alpha}$ and the population size for conditional sampling. Considering the heavy-tail behaviour of $X = g(\mathbf{U})$, the estimator has a finite variance as soon as $a > 1 + 1/N = 1.05$. One the one hand we know here that the tail-index of X is equal to $1/0.8 = 1.25$; on the other hand it is easy to check this condition afterwards by estimating the slope on the plot $\log X$ against N_{iter}/N as in Figure 2.

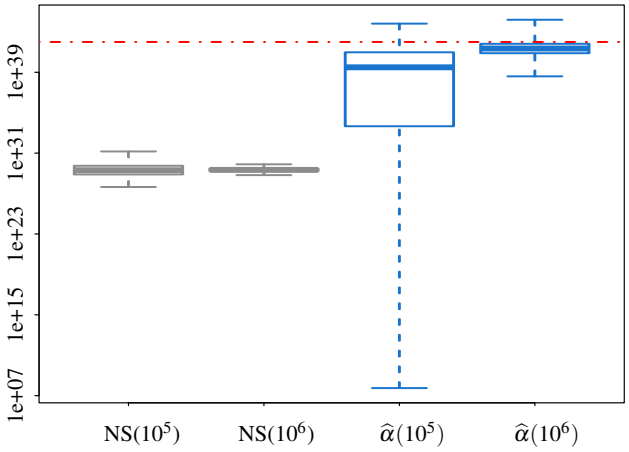

Fig. 5: Estimation of $\mathrm{E}\left[g_{\text{ht}}(\mathbf{U})\right]$ with $\mathbf{U} \sim \mathscr{U}\left(-[\frac{1}{2},\frac{1}{2}]^d\right)$, $d = 20$. (NS): nested sampling stopped after $100N$ iterations; $\widehat{\alpha}$: estimator of Section 3.4 with β_{app} (Eq. (13)) and $N = 20$. 10^5 and 10^6 are the computational budgets used. The (red) dot-dashed line is the theoretical value of m.

It is visible on Figure 5 that nested sampling did not go far enough and misses an important part of the mass: $\mathrm{E}\left[\mathrm{NS}(10^5)\right] = 5.32 \times 10^{29}$ and $\mathrm{E}\left[\mathrm{NS}(10^6)\right] = 2.41 \times 10^{29}$ while the reference value is 1.08×10^{42}. On the other hand, $\widehat{\alpha}$ is unbiased (estimated means are 6.43×10^{41} and 1.52×10^{42}). However, it does not seem to be approximately Gaussian yet. Indeed $\widehat{Z}$ can be relatively heavy-tailed (McLeish 2011) and a consequent computational budget may be required for $\widehat{\alpha}$ to effectively become normally distributed.

6 Conclusion

Nested Sampling has been proposed as a method for estimating the evidence in a Bayesian framework and applied with success in a great variety of areas like astronomy and cosmology. Since its introduction, a lot of work has been done to clarify its convergence properties (*e.g.* Evans 2007; Chopin and Robert 2010; Keeton 2011) and to handle the issue of conditional sampling (*e.g.* Mukherjee et al 2006; Brewer et al 2011; Martiniani et al 2014). However nested sampling termination remains an open issue and a matter of user judgement (Skilling 2006, Section 7).

Linking nested sampling with recent results in rare event simulation, this paper extends it to the estimation of the mean of any real-valued random variable (being bounded or not) and goes on step further by giving the optimal nested sampling weights and proving that 1) an idealised nested sampling with slightly modified weights and an infinite number of iterations is unbiased; 2) its variance is always lower than the classical Monte Carlo estimator one's; and 3) the random variable of interest does not need to have a finite second-order moment to produce an estimator with finite variance. This latter property makes nested sampling especially relevant for heavy-tailed random variables as developed Section 4.

Furthermore, we also present two ways of implementing a practical unbiased estimator with an a.s. finite number of terms, resolving the issue of choosing an arbitrary stopping criterion. The first estimator can be used exactly as usual nested sampling and preserves unbiasedness while only doubling the variance of the ideal estimator (infinite number of terms). The second one can be used with a predetermined fixed computational budget and supports a Central Limit Theorem. Practically speaking, they both enable parallel implementation (unlike usual adaptive nested sampling strategies) and do not depend on the random variable of interest.

As for any nested sampling implementations, they require to be able to generate samples according to conditional laws and theoretical results are derived with this hypothesis. In some cases, exact conditional sampling may be possible. When the random variable of interest is the output of a computer code, Markov Chain drawing like Metropolis-Hastings algorithm can overcome this issue. If only *iid* samples are available, further work has to be done to explicit the link between the increasing random walk presented in Section 2.1 and, for example, Pareto-type distributions.

Acknowledgements The author would like to thank his advisors Josselin Garnier (University Paris Diderot) and Gilles Defaux (Commissariat à l'Energie Atomique et aux Energies Alternatives) for their advices and suggestions as well as the reviewers for their very relevant comments which helped improving the manuscript. This work was partially supported by ANR project Chorus.

Appendix

Proof of Proposition 2 one has:

$$\mathrm{E}\left[\widehat{m}\right] = \int_0^\infty \mathrm{E}\left[\left(1 - \frac{1}{N}\right)^{M_x}\right] \mathrm{d}x = \int_0^\infty p_x \mathrm{d}x.$$

For the variance, one uses the fact that, for $x > x'$, $M_x - M_{x'}$ and $M_{x'}$ are independent to expand $\mathrm{E}\left[\widehat{m}^2\right]$:

$$\begin{aligned}\mathrm{E}\left[\widehat{m}^2\right] &= 2\int_0^\infty \int_0^x \mathrm{E}\left[\left(1-\frac{1}{N}\right)^{M_x+M_{x'}}\right] \mathrm{d}x'\mathrm{d}x \\ &= \int_0^\infty \int_0^x \mathrm{E}\left[\left(1-\frac{1}{N}\right)^{M_x-M_{x'}}\left(1-\frac{1}{N}\right)^{2M_{x'}}\right] \mathrm{d}x'\mathrm{d}x.\end{aligned}$$

Furthermore renewal property of a Poisson process gives $M_x - M_{x'} \sim \mathscr{P}(-\log(p_x/p_{x'}))$. Eventually one can conclude using the results of Proposition 1. □

Proof of Proposition 3 Starting from the expression of the variance found in Proposition 2:

$$\mathrm{var}\left[\widehat{m}\right] = 2\int_0^\infty p_x \int_0^x p_{x'}^{1-1/N}\mathrm{d}x'\mathrm{d}x - \mathrm{E}\left[X\right]^2,$$

we make use of Hölder's inequality:

$$\begin{aligned}&\int_0^x p_{x'}^{1-1/N} dx' \\ &\leq \left(\int_0^x dx'\right)^{1/N}\left(\int_0^x p_{x'}dx'\right)^{1-1/N} \\ &\leq x^{1/N}\left(\int_0^\infty p_{x'}\mathrm{d}x'\right)^{1-1/N} \\ &\leq x^{1/N}\,\mathrm{E}\left[X\right]^{1-1/N}.\end{aligned}$$

And therefore:

$$\mathrm{var}\left[\widehat{m}\right] \leq \frac{2}{1+1/N}\,\mathrm{E}\left[X\right]^{1-1/N}\mathrm{E}\left[X^{1+1/N}\right].$$

Using Hölder's inequality again, one gets:

$$\mathrm{var}\left[\widehat{m}\right] \leq \frac{2}{1+1/N}\,\mathrm{E}\left[X^{1+1/N}\right]^{\frac{2}{1+1/N}}.$$

□

Proof of Proposition 4 On the one hand one has:

$$N\,\mathrm{var}\left[\widehat{m}_{MC}\right] + m^2 = 2\int_0^\infty x p_x \mathrm{d}x,$$

and on the other hand one can write:

$$N\,\mathrm{var}\left[\widehat{m}\right] + m^2 = 2\int_0^\infty p_x \int_0^x p_{x'}\left[N(p_{x'}^{-1/N}-1)+1\right]dx'dx.$$

Considering $f : p \mapsto p\left[N(p^{-1/N}-1)+1\right]$, we have $f(1) = 1$ and:

$$f'(p) = (N-1)(p^{-1/N}-1) \geq 0\,, \forall p \in [0,1].$$

Thus: $\forall p \in [0,1], f(p) \leq 1$. Therefore

$$N\,\mathrm{var}\left[\widehat{m}\right] + m^2 \leq 2\int_0^\infty x p_x \mathrm{d}x$$

which shows that $\mathrm{var}\left[\widehat{m}\right] \leq \mathrm{var}\left[\widehat{m}_{MC}\right]$. □

Proof of Proposition 5 Starting with the last formulation in (10) for $\widehat{Z}$, one uses the fact that T and $(X_i)_i$ are independent. Finally, (4) and Proposition 2 let conclude: $\mathrm{E}\left[\widehat{Z}\right] = m$.

For the second-order moment, we use the fact that $\widehat{Z}$, like $\widehat{m}$, can be written with an integral:

$$\widehat{Z} = \int_0^\infty \left(1-\frac{1}{N}\right)^{M_x} \frac{\mathbb{1}_{T\geq M_x}}{\mathrm{P}\left[T \geq M_x\right]}\mathrm{d}x$$

and apply the same reasoning as for $\mathrm{E}\left[\widehat{m}^2\right]$: given $x > x'$, the random variables $M_x - M_{x'}$, $M_{x'}$ and T are independent, which brings:

$$\begin{aligned}&\mathrm{E}\left[\left(1-\frac{1}{N}\right)^{M_x+M_{x'}} \frac{\mathbb{1}_{T\geq M_x}}{\mathrm{P}\left[T\geq M_x\right]}\frac{\mathbb{1}_{T\geq M_{x'}}}{\mathrm{P}\left[T\geq M_{x'}\right]}\right] \\ &= \mathrm{E}\left[\left(1-\frac{1}{N}\right)^{M_x-M_{x'}}\left(1-\frac{1}{N}\right)^{2M_{x'}}\beta_{M_{x'}}^{-1}\frac{\mathbb{1}_{T\geq M_x}}{\mathrm{P}\left[T\geq M_x\right]}\right] \\ &= \mathrm{E}\left[\left(1-\frac{1}{N}\right)^{M_x-M_{x'}}\left(1-\frac{1}{N}\right)^{2M_{x'}}\beta_{M_{x'}}^{-1}\right] \\ &= \frac{p_x}{p_{x'}}\sum_{i=0}^\infty e^{N\log p_{x'}}\frac{\left[-N\log p_{x'}(1-1/N)^2\right]^i}{i!}\beta_i^{-1} \\ &= \sum_{i=0}^\infty p_x p_{x'}^{N-1}\frac{\left[-N\log p_{x'}(1-1/N)^2\right]^i}{i!}\beta_i^{-1}.\end{aligned}$$

Then using this equality in $\mathrm{E}\left[\widehat{Z}^2\right]$ gives the solution. □

Proof of Lemma 1 Let $\varepsilon > 0$ be such that $\mathrm{E}\left[X^{1+\varepsilon}\right] < \infty$, $N \in \mathbb{N} \mid N > 1/\varepsilon$ and $i \geq 0$. We further extend the definition of $\mathrm{var}\left[\widehat{m}\right]$ given in Proposition 1, Eq. (8) for any $N \in \mathbb{R}$. Proof of Proposition 3 is based on Hôlder's inequality and still holds in this case, and so for Corollary 1. Hence, according to Corollary 1: $\exists N' \in \mathbb{R}$ such that $N' < N$ and $\mathrm{var}\left[\widehat{m}\right](N') < \infty$. Furthermore, given x and x' one can write:

$$p_x p_{x'}^{N-1}(-\log p_{x'})^i = p_x p_{x'}^{1-1/N'} p_{x'}^{N+1/N'-2}(-\log p_{x'})^i.$$

Moreover the function $p : (0,1) \mapsto p^{N+1/N'-2}(-\log p)^i$ is bounded above by $e^{-i}i^i(N+1/N'-2)^{-i}$. Using the Stirling lower bound $i \geq i^i e^{-i}\sqrt{2\pi i}$ we can write:

$$p_x p_{x'}^{N-1}(-\log p_{x'})^i \leq p_x p_{x'}^{1-1/N'}\frac{i!}{\sqrt{2\pi i}(N+1/N'-2)^i}.$$

Finally, this inequality brings:

$$q_{i,N} \leq \mathrm{var}\left[\widehat{m}\right](N')\left(\frac{N(1-1/N)^2}{N+1/N'-2}\right)^i\frac{1}{\sqrt{2\pi i}}$$

and $(N+1/N-2)/(N+1/N'-2) < 1$, which concludes the first part of the proof.

Let us now assume that X has a density f_X. One has:

$$q_{i,N} = 2\int_0^\infty \int_0^x p_x p_{x'}^{N-1} \frac{\left[-N\log p_{x'}(1-1/N)^2\right]^i}{i!}\mathrm{d}x'\mathrm{d}x.$$

Denote x_L the left end point of X (remember that X is non-negative valued so $x_L \geq 0$). Then:

$$q_{i,N} \geq 2\int_{x_L}^\infty \int_{x_L}^x p_x p_{x'}^{N-1} \frac{\left[-N\log p_{x'}(1-1/N)^2\right]^i}{i!}\mathrm{d}x'\mathrm{d}x.$$

We then consider the change of variable $u = -\log p_x$ and $u' = -\log p_{x'}$; for all $i \geq 1$ one has:

$$\begin{aligned}
q_{i,N} &\geq \frac{2}{\|f_X\|_\infty^2}\left(1-\frac{1}{N}\right)^{2i}\int_0^\infty e^{-2u}\int_0^u \frac{e^{-Nu'}(Nu')^i}{i!}\mathrm{d}u'\mathrm{d}u \\
&\geq \frac{2}{\|f_X\|_\infty^2}\left(1-\frac{1}{N}\right)^{2i}\int_0^\infty e^{-2u}\frac{1}{N}\sum_{k=i+1}^\infty \frac{e^{-Nu}(Nu)^k}{k!}\mathrm{d}u \\
&\geq \frac{2}{\|f_X\|_\infty^2}\frac{1}{N(N+2)}\left(1-\frac{1}{N}\right)^{2i}\sum_{k=i+1}^\infty \left(\frac{N}{N+2}\right)^k \\
q_{i,N} &\geq \frac{1}{(N+2)\|f_X\|_\infty^2}\left[\frac{N}{N+2}\left(1-\frac{1}{N}\right)^2\right]^i.
\end{aligned}$$

□

Proof of Proposition 6 Let $\alpha > 0$ be such that $(1-1/N) = e^{-\alpha}$. The argument is the same one as in Proposition 5. One has:

$$\begin{aligned}
\mathrm{E}\left[\widehat{Z}^2\right] &= 2\int_0^\infty\int_0^x \mathrm{E}\left[e^{-\alpha(M_x-M_{x'})}e^{(\beta-2\alpha)M'_x}\right]\mathrm{d}x'\mathrm{d}x \\
&= 2\int_0^\infty\int_0^x p_x p_{x'}^{1-\frac{1}{\gamma(\beta,N)}}\mathrm{d}x'\mathrm{d}x
\end{aligned}$$

with:

$$\begin{aligned}
\frac{N}{\gamma(\beta,N)} &= 2N - N^2 + e^\beta(N-1)^2 \\
&= 1 + (N-1)^2(e^\beta - 1).
\end{aligned}$$

□

Proof of Corollary 2 Noticing that for any $N \geq 2$, one has $\gamma(\beta_{\mathrm{app}}(N),N) = (N+1)/2$ gives the first equality. Then, since $\mathrm{var}[\widehat{m}]$ typically scales with $1/N$ (usual results on nested sampling) gives the approximation.

Proof of Proposition 7 If $T = 0$ then no other simulation is done other than the first element of each Markov chain, *i.e.* N simulations are done. Then each step requires the simulation of the next stopping time, *i.e.* one simulation. Finally, this brings $\tau = N + T$.

Proof of Corollary 3 Note that $\mathrm{var}[\widehat{m}] = \sum_{i=0}^\infty q_{i,N} - m^2$. Hence, one has $\mathrm{var}\left[\widehat{Z}\right] > \mathrm{var}[\widehat{m}]$ because $\mathrm{var}\left[\widehat{Z}\right] = \mathrm{var}[\widehat{m}] \Leftrightarrow \forall i \in \mathbb{N},\ \beta_i = \mathrm{P}[T \geq i] = 1$ and $\mathrm{E}[\tau] > N$ because $\mathrm{E}[\tau] = N \Leftrightarrow \mathrm{E}[T] = 0$ while $\forall i \in \mathbb{N},\ \mathrm{P}[T \geq i] > 0$. Furthermore, the power series expansion of the exponential function and the dominated convergence theorem let us rewrite $\mathrm{var}[\widehat{m}]$:

$$\begin{aligned}
\mathrm{var}[\widehat{m}] &= \sum_{i=1}^\infty 2\int_0^\infty\int_{x'}^\infty p_x p_{x'}\frac{(-\log p_{x'})^i}{N^i i!}\mathrm{d}x\mathrm{d}x' \\
\mathrm{var}[\widehat{m}] &= \sum_{i=1}^\infty q_{i,2}\left(\frac{2}{N}\right)^i
\end{aligned}$$

which brings: $\mathrm{var}[\widehat{m}] = q_{1,2}\cdot 2/N + O\left(1/N^2\right)$. All together, these inequalities complete the proof. □

Proof of Corollary 4 Denote $B = 1/(e^\beta - 1)$; one has:

$$\frac{N+B}{\gamma(B,N)} = N + \frac{B}{N} + \frac{N^2}{B} - 1 - \frac{2N}{B} + \frac{1}{B} + \frac{1}{N}.$$

With $\beta = \Theta(1/N^{1+\varepsilon})$, $\varepsilon \geq 0$, one has $B \sim 1/\beta \sim N^{1+\varepsilon}$. Finally, this gives:

$$\frac{N+B}{\gamma(B,N)} \sim N + N^\varepsilon + N^{1-\varepsilon} + O(1),$$

which concludes the proof. □

Proof of Proposition 8 First one shows that i_0 is well determined. The sequence $(\Delta_i)_i$ defined by:

$$\forall i \in \mathbb{N}, \Delta_i = \sum_{j=0}^i q_{j,N} - m^2 - (N+i)q_{(i+1),N}$$

is increasing:

$$\begin{aligned}
\Delta_{i+1} - \Delta_i &= q_{(i+1),N} - (N+i+1)q_{(i+2),N} + (N+i)q_{(i+1),N} \\
&= (N+i+1)(q_{(i+1),N} - q_{(i+2),N}) > 0.
\end{aligned}$$

Furthermore $q_0 - m^2 = 2\int_0^\infty\int_{x'}^\infty p_x p_{x'}\left(p_{x'}^{N-2} - 1\right)\mathrm{d}x\mathrm{d}x' \leq 0 < Nq_{1,N}$, so $\Delta_0 < 0$, and $\Delta_i \to \mathrm{var}[\widehat{m}]$ when $i \to \infty$ because $(q_{i,N})_i$ decreases at exponential rate. So there exists $i_0 \in \mathbb{N} \mid \Delta_{i_0-1} \leq 0$ and $\Delta_{i_0} > 0$.

Let us now consider the auxiliary problem:

$$\underset{\substack{(\beta_i)_{i\geq 1}\\ \beta_i > 0}}{\mathrm{argmin}} \left(\beta + \sum_{i=1}^\infty \beta_i\right)\left(q + \sum_{i=1}^\infty q_{i,N}\beta_i^{-1}\right)$$

with $\beta > 0$ and $q \in \mathbb{R}$. We show that it has a solution if and only if $q > 0$. Let $i \geq 1$, cancelling the partial derivatives brings:

$$\forall i \geq 1,\ 0 = \left(q + \sum_{j=1}^\infty q_j\beta_j^{-1}\right) + \left(\beta + \sum_{j=1}^\infty \beta_j\right)\frac{-q_{i,N}}{\beta_i^2}.$$

Then the solution should be of the form: $\forall i \in [\![1,\infty)\!), \beta_i = c_0\sqrt{q_i}$ for some $c_0 > 0$. Solving now the problem with c_0, the derivative writes $q - \beta/c_0^2$. If $q \leq 0$ then it is strictly decreasing and there is no global minimiser. On the contrary, $q > 0$ brings $c_0 = \sqrt{\beta/q}$ and $\forall i \geq 1, \beta_i = c_0\sqrt{q_i}$.

Thus, in our context with the constraint $\forall i \in \mathbb{N}, \beta_i \leq 1$, this means that solving the optimisation problem will set iteratively $\beta_i = 1$ until the minimiser is feasible, *i.e.* until $i_0 \stackrel{\text{def}}{=} \min\{i \in \mathbb{N} \mid \sum_{j=0}^{i} q_{j,N} - m^2 > (N+i)q_{(i+1),N}\}$. Then the solution will be given by:

$$\forall i \in [\![1, i_0]\!], \beta_i = 1$$
$$\forall i > i_0, \beta_i = \frac{\sqrt{q_{i,N}}}{\sqrt{\frac{1}{N+i_0}\sum_{j=0}^{i_0}(q_{j,N} - m^2)}}.$$

□

Proof of Corollary 5 By definition of i_0, one has:

$$(N+i_0)q_{i_0+1} < \sum_{j=0}^{i_0} q_j - m^2 \leq (N+i_0-1)q_{i_0} + q_{i_0}$$

which concludes the proof. □

Proof of Proposition 9 Denote:

$$Q_N(\beta) = \left(N + \frac{1}{e^\beta - 1}\right)\left(\sum_{i=0}^{\infty} q_{i,2}(2/\gamma)^i - m^2\right)$$

the quantity one seeks to minimise.

First, we show that for any fixed N, there exists a global minimiser of $Q_N(\beta)$. One has $Q_N(\beta) \to \infty$ when $\beta \to 0$ and $\gamma(\beta, N) \to 0$ when $\beta \to \infty$. Hence, either $\exists \beta_\infty \in (0,\infty]$ such that:

$$\begin{cases} Q_N(\beta) \xrightarrow[\beta \nearrow \beta_\infty]{} \infty \\ Q_N(\beta) < \infty \qquad \forall \beta < \beta_\infty. \end{cases}$$

Then Q_N is continuous on $(0, \beta_\infty)$ with infinite limits on 0 and β_∞, so it reaches its minimum on $(0, \beta_\infty)$; or $\exists \beta_\infty \in (0,\infty)$ such that:

$$\begin{cases} Q_N(\beta) < \infty \quad \forall \beta \in (0, \beta_\infty] \\ Q_N(\beta) = \infty \quad \forall \beta > \beta_\infty. \end{cases}$$

Since Q_N is continuous on β_∞^- by Monotone Convergence Theorem, Q_N reaches its minimum on $(0, \beta_\infty]$.

Let $\beta_{\text{opt}}(N) > 0$ be such that $\inf_\beta Q_N(\beta) = Q_N(\beta_{\text{opt}})$. We now show that there exists an optimal N. It is sufficient to show $Q_N(\beta_{\text{opt}}) \to \infty$ when $N \to \infty$. Denote $B = 1/(e^\beta - 1)$; one has:

$$\frac{1}{\gamma(B,N)} = \frac{1}{N} + \frac{N}{B} - \frac{2}{B} + \frac{1}{NB}.$$

Hence, depending on the growth rate of B when $N \to \infty$, one would have:

$$B = O(N), \frac{1}{\gamma} \sim \frac{N}{B} \Rightarrow \inf_\beta Q_N(\beta) \xrightarrow[N\to\infty]{} \infty$$
$$N = o(B), \frac{1}{\gamma} \sim \frac{1}{N} \text{ or } \frac{N}{B} \Rightarrow \inf_\beta Q_N(\beta) \sim \frac{B}{N} \text{ or } N$$
$$\Rightarrow \inf_\beta Q_N(\beta) \xrightarrow[N\to\infty]{} \infty.$$

Then in any cases $Q_N(\beta_{\text{opt}}) \to \infty$ when $N \to \infty$, which means that there exists $N_{\text{opt}} \in \mathbb{N} \mid Q_{N_{\text{opt}}}(\beta_{\text{opt}}) = \inf_N Q_N(\beta_{\text{opt}})$.

We now show the relationship between β_{opt} and N_{opt}: the partial derivatives of $\mathrm{E}[\tau] \cdot \mathrm{var}\left[\widehat{Z}\right]$ against B and N write:

$$\begin{cases} \dfrac{\partial\left(\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right]\right)}{\partial B} = \mathrm{var}\left[\widehat{Z}\right] + \mathrm{E}[\tau]\dfrac{\partial \mathrm{var}\left[\widehat{Z}\right]}{\partial \gamma}\dfrac{\partial \gamma}{\partial B} \\ \dfrac{\partial\left(\mathrm{E}[\tau]\cdot\mathrm{var}\left[\widehat{Z}\right]\right)}{\partial N} = \mathrm{var}\left[\widehat{Z}\right] + \mathrm{E}[\tau]\dfrac{\partial \mathrm{var}\left[\widehat{Z}\right]}{\partial \gamma}\dfrac{\partial \gamma}{\partial N}. \end{cases}$$

At point $(\beta_{\text{opt}}, N_{\text{opt}})$, both equations are cancelled, which gives:

$$\frac{\partial \gamma}{\partial N}(B_{\text{opt}}, N_{\text{opt}}) = \frac{\partial \gamma}{\partial B}(B_{\text{opt}}, N_{\text{opt}}).$$

Recalling $\gamma(B,N) = NB/(B + (N-1)^2)$, this gives the equation: $B_{\text{opt}}^2 - (N_{\text{opt}}^2 - 1)B_{\text{opt}} - N_{\text{opt}}(N_{\text{opt}} - 1)^2 = 0$. One can solve it in B_{opt} and keep the positive root, which gives the solution.

Proof of Proposition 10 For the first equality:

$$\mathrm{E}[X] = \int_0^\infty p_x \mathrm{d}x = \frac{a}{a-1}$$
$$\mathrm{var}[\widehat{m}_{MC}] = \frac{1}{N}\left(\mathrm{E}\left[X^2\right] - \mathrm{E}[X]^2\right) = \frac{a}{N(a-2)(a-1)^2}$$
$$= \frac{m(m-1)^2}{(2-m)N};$$

for the second one:

$$\mathrm{E}\left[\widehat{m}^2\right] = 2\int_0^\infty \int_0^x p_x p_{x'}^{1-1/N} \mathrm{d}x' \mathrm{d}x$$
$$= 2\int_0^1 \int_0^x \cdots + 2\int_1^\infty \int_0^1 \cdots + 2\int_1^\infty \int_1^x \cdots$$
$$= 1 + \frac{2}{a-1} + \frac{2}{(a-1)(2(a-1) - a/N)}$$
$$\mathrm{var}[\widehat{m}] = \frac{a}{N(a-1)^2(2(a-1) - a/N)};$$

and for the third one:

$$\mathrm{var}[\widehat{m}_{IS}] = \frac{1}{N}\left[\int_1^\infty x^2 \frac{a^2}{b} x^{-2a+b-1} \mathrm{d}x - \frac{a^2}{(a-1)^2}\right]$$
$$\mathrm{var}[\widehat{m}_{IS}] = \frac{a^2}{N(a-1)^2}\left(\frac{1}{B(2-B)} - 1\right)$$

with $B = b/(a-1)$. □

Proof of Proposition 11 Let $i \geq 0$, one has:

$$\int_1^\infty \int_{x'}^\infty p_x p_{x'}^{N-1} \frac{\left[-N \log p_{x'}(1-1/N)^2\right]^i}{i!} \mathrm{d}x\mathrm{d}x' =$$
$$\frac{\left[aN(1-1/N)^2\right]^i}{i!} \int_1^\infty \int_{x'}^\infty x^{-a} x'^{-a(N-1)} (\log x')^i \mathrm{d}x\mathrm{d}x'$$
$$= \frac{\left[aN(1-1/N)^2\right]^i}{(a-1)i!} \int_1^\infty x'^{1-aN} (\log x')^i \mathrm{d}x'$$
$$= \frac{\left[aN(1-1/N)^2\right]^i}{(a-1)i!} \frac{\Gamma(i+1)}{(aN-2)^{i+1}}$$
$$= \frac{1}{(a-1)(aN-2)} \left[\frac{aN}{aN-2}\left(1-\frac{1}{N}\right)^2\right]^i$$

with Γ standing here for the Gamma function. Furthermore:

$$\int_0^1 \int_{x'}^\infty p_x p_{x'}^{N-1} \frac{\left[-N \log p_{x'}(1-1/N)^2\right]^i}{i!} \mathrm{d}x\mathrm{d}x' = \frac{\mathbb{1}_{i=0}(a+1)}{2(a-1)}.$$

$(q_{i,N})_i$ is decreasing iff:

$$\frac{aN}{aN-2}\left(1-\frac{1}{N}\right)^2 < 1 \Leftrightarrow 1 < a\left(1-\frac{1}{2N}\right)$$

which is indeed the condition for the finiteness of $\mathrm{var}[\widehat{m}]$ already stated in Proposition 10. □

Proof of Proposition 12 The problem can be rewritten:

$$\min\left\{i \geq 1 \mid \frac{1}{1-\beta} - \frac{aN-2}{2(a-1)} > \beta^{i+1}\left(N+i+\frac{1}{1-\beta}\right)\right\}.$$

Furthermore one has:

$$\frac{1}{1-\beta} = \frac{Nm}{2} + \frac{(a-2)^2}{4(a-1)^2} + o(1)$$

which brings that the left hand term is equal to $(m/2)^2 + o(1)$. Writing $i = N(k_0 + k_1 \log N + k_2 \log\log N)$ brings:

$$\beta^{i+1} = e^{-\frac{2k_0}{m}} N^{-\frac{2k_1}{m}} (\log N)^{-\frac{2k_2}{m}} (1+o(1)).$$

Hence one has to choose k_0, k_1 and k_2 such that the right hand term also equals $(m/2)^2 + o(1)$, which gives the solution. □

Proof of Corollary 6 Using the asymptotic expansion of i_0 one finds $q_{i_0} \sim (N^2 \log N)^{-1}(m-1)^2$. Furthermore, one has $\mathrm{E}[\tau] \sim i_0$. Finally, the use of $\mathrm{E}[\tau] \cdot \mathrm{var}\left[\widehat{Z}\right] \sim q_{i_0} \mathrm{E}[\tau]^2$ gives the result. □

Proof of Proposition 13 One gets the expression of the variance directly from Section 2.2 with $\gamma(N,\beta)$ instead of N. Then, denoting $B = 1/(e^\beta - 1)$, one solves the problem:

$$\frac{\partial}{\partial B}\left((N+B)\left(\frac{a}{2(a-1)\gamma - a}\right)\right) = 0.$$

□

References

Au SK, Beck JL (2001) Estimation of small failure probabilities in high dimensions by subset simulation. Probabilistic Engineering Mechanics 16(4):263–277

Beirlant J, Caeiro F, Gomes MI (2012) An overview and open research topics in statistics of univariate extremes. REVSTAT-Statistical Journal 10(1):1–31

Bernardo JM, Bayarri M, Berger JO, Dawid AP, Heckerman D (2011) Bayesian Statistics 9. Oxford University Press

Brewer BJ, Pártay LB, Csányi G (2011) Diffusive nested sampling. Statistics and Computing 21(4):649–656

Cérou F, Guyader A (2007) Adaptive multilevel splitting for rare event analysis. Stochastic Analysis and Applications 25(2):417–443

Cérou F, Del Moral P, Furon T, Guyader A, et al (2009) Rare event simulation for a static distribution. Proceedings of RESIM 2008 URL http://www.irisa.fr/aspi/fcerou/Resim_Cerou_et_al.pdf

Cérou F, Del Moral P, Furon T, Guyader A (2012) Sequential Monte Carlo for rare event estimation. Statistics and Computing 22(3):795–808

Chopin N, Robert CP (2010) Properties of nested sampling. Biometrika p asq021

Corless RM, Gonnet GH, Hare DE, Jeffrey DJ, Knuth DE (1996) On the lambertw function. Advances in Computational mathematics 5(1):329–359

Embrechts P, Klüppelberg C, Mikosch T (1997) Modelling extremal events: for insurance and finance, vol 33. Springer

Evans M (2007) Discussion of nested sampling for bayesian computations by john skilling. Bayesian Statistics 8:491–524

Garvels MJJ (2000) The splitting method in rare event simulation. Universiteit Twente

Giles MB (2008) Multilevel monte carlo path simulation. Operations Research 56(3):607–617

Glynn PW, Iglehart DL (1989) Importance sampling for stochastic simulations. Management Science 35(11):1367–1392

Glynn PW, Whitt W (1992) The asymptotic efficiency of simulation estimators. Operations Research 40(3):505–520

Guyader A, Hengartner N, Matzner-Løber E (2011) Simulation and estimation of extreme quantiles and extreme probabilities. Applied Mathematics & Optimization 64(2):171–196

Hill JB (2013) Robust estimation for average treatment effects. Available at SSRN 2260573 URL http://dx.doi.org/10.2139/ssrn.2260573

Huber M, Schott S, et al (2011) Using tpa for bayesian inference. Bayesian Statistics 9 9:257

Huber M, Schott S, et al (2014) Random construction of interpolating sets for high-dimensional integration. Journal of Applied Probability 51(1):92–105

Johansson J (2003) Estimating the mean of heavy-tailed distributions. Extremes 6(2):91–109

Kahn H, Harris TE (1951) Estimation of particle transmission by random sampling. National Bureau of Standards applied mathematics series 12:27–30

Keeton CR (2011) On statistical uncertainty in nested sampling. Monthly Notices of the Royal Astronomical Society 414(2):1418–1426

Martiniani S, Stevenson JD, Wales DJ, Frenkel D (2014) Superposition enhanced nested sampling. Physical Review X 4(3):031,034

McLeish D (2011) A general method for debiasing a monte carlo estimator. Monte Carlo Methods and Applications

Mukherjee P, Parkinson D, Liddle AR (2006) A nested sampling algorithm for cosmological model selection. The Astrophysical Journal Letters 638(2):L51

Necir A, Rassoul A, Zitikis R (2010) Estimating the conditional tail expectation in the case of heavy-tailed losses. Journal of Probability and Statistics 2010, URL http://dx.doi.org/10.1155/2010/596839

Peng L (2001) Estimating the mean of a heavy tailed distribution. Statistics & Probability Letters 52(3):255–264

Propp JG, Wilson DB (1996) Exact sampling with coupled markov chains and applications to statistical mechanics. Random structures and Algorithms 9(1-2):223–252

Rhee Ch, Glynn PW (2013) Unbiased estimation with square root convergence for sde models. Submitted for publication URL http://rhee.gatech.edu/papers/RheeGlynn13a.pdf

Robert CP, Casella G (2004) Monte Carlo statistical methods. Springer

Roberts G (2011) Comments on "Using TPA for Bayesian inference" by Huber, M. and Schott, S. In: Bernardo JM, Bayarri MJ, Berger JO, Dawid AP, Heckerman D, M SAF, West M (eds) Bayesian Statistics 9, Oxford University Press, pp 257–282

Simonnet E (2014) Combinatorial analysis of the adaptive last particle method. Statistics and Computing pp 1–20

Skilling J (2006) Nested sampling for general bayesian computation. Bayesian Analysis 1(4):833–859

Vergé C, Dubarry C, Del Moral P, Moulines E (2013) On parallel implementation of sequential monte carlo methods: the island particle model. Statistics and Computing pp 1–18

Walter C (2015) Moving particles: A parallel optimal multilevel splitting method with application in quantiles estimation and meta-model based algorithms. Structural Safety 55(0):10 – 25